arXiv:1410.5000v2 [cs.LO] 17 Feb 2015

On the complexity of Linearizability

Jad Hamza

LIAFA, Université Paris Diderot

Abstract. It was shown in Alur et al. [1] that the problem of verifying finite concurrent systems through Linearizability is in EXPSPACE. However, there was still a complexity gap between the easy to obtain PSPACE lower bound and the EXPSPACE upper bound. We show in this paper that Linearizability is EXPSPACE-complete.

1 Introduction

Linearizability [8] is the standard consistency criterion for concurrent data-structures. Filipovic et al. [5] proved that checking that a library L is linearizable with respect to a specification S is equivalent to observational refinement. Formally, as long as linearizability holds, any multi-threaded program using the specification S as a library can safely replace it by L, without adding any unwanted behaviors.

Many practical tools [3, 12, 4, 13, 11] for checking linearizability or detecting linearizability violations exist, and here is short summary of the work done on the complexity.

Checking that a single execution is linearizable is already an NP-complete problem[7]. Moreover, Alur et al. [1] showed that the problem of checking Linearizability for a finite concurrent libraries used by a finite number of threads is in EXPSPACE when the specification is a regular language. The best known lower bound is PSPACE-hardness, obtained from a simple reduction of the reachability problem for finite concurrent programs [1], leaving a large complexity gap.

This result was refined in Bouajjani et al. [2] where it was shown that a simpler variant of Linearizability – called Static Linearizability, or Linearizability with fixed linearization points – is PSPACE-complete for the same class of libraries.

Furthermore, Linearizability is undecidable when the number of threads is unbounded [2]. Tools used for detecting linearizability violations often start by underapproximating the set of executions by bounding the number of threads. It is thus necessary to develop a better understanding of Linearizability for a bounded number of threads.

We prove that Linearizability is EXPSPACE-complete, showing that there is an inherent difficulty to the problem. We introduce for this a new problem on regular languages, called Letter Insertion. This problem can be reduced in polynomial time to Linearizability.

We then show that Letter Insertion is EXPSPACE-hard, closing the complexity gap for Linearizability. Our proof is similar to the proofs of EXPSPACE-

hardness for the problems of inclusion of extended regular expressions with intersection operator, or interleaving operator, given in Hunt [9], Fürer [6] and Mayer and Stockmeyer [10]. They all use a similar encoding of runs of Turing machines as words, and using the problem at hand, Letter Insertion in this case, to recognize erroneous runs.

To summarize, our two contributions are:

- finding the Letter Insertion problem, a problem equivalent to Linearizability, but which has a very simple formulation in terms of regular automata,
- using this problem to show EXPSPACE-hardess of Linearizability.

We recall in Section 2 the definition of Linearizability, and we introduce the Letter Insertion problem. We show in Section 3 that Letter Insertion can be reduced in polynomial time to Linearizability. And finally, we show in Section 4, that Letter Insertion is EXPSPACE-hard, which is the most technical part of the paper. When combined, Sections 3 and 4 show that Linearizability is EXPSPACE-hard.

2 Definitions

2.1 Libraries

In the usual sense, a *library* is a collection of *methods* that can be called by other programs. We start by giving our formalism for methods, and define libraries as sets of methods.

In order to simplify the presentation, and since they do not affect our EXPSPACE-hardness reduction, we will use a number of restrictions on the methods. First, we will define the methods without return values and parameters. Second, each instruction of a method can either read or write to the shared memory, but we don't formalize atomic compare and set operations. Finally, we limit ourselves to a unique *shared variable*.

Let $\mathbb{D}$ be a finite set used as the *domain* for the shared variable and let $d_0 \in \mathbb{D}$ be a special value considered as initial.

A *method* is a tuple (Q, δ, q_0, q_f) where

- Q is the set of states,
- $\delta \subseteq Q \times \{\mathsf{read}, \mathsf{write}\} \times \mathbb{D} \times Q$
- $q_0 \in Q$ is the initial state (in which the method is called)
- $q_f \in Q$ is the final state (in which the method can return)

One point which might be considered unsual in our formalism is that a read instruction *guesses* the value that it is going to read. In usual programming languages, this can be understood as first reading a variable, and then having an assume statement to constrain the value of the read variable. This formalism choice is a presentation choice, and has no effect on the complexity of the problem.

As hinted previously, a *library* $Lib = \{M_1, \ldots, M_m\}$ is a set of methods. For every $j \in \{1, \ldots, m\}$, let $(Q^j, \delta^j, q_0^j, q_f^j)$ be the tuple corresponding to M_j. We define Q to be the (disjoint) union of all Q^j.

Let k be an integer representing the number of *threads* using Lib. Threads run concurrently and call the methods of Lib arbitrarily. The system composed of k threads calling arbitrarily the methods of Lib is called Lib^k.

Formally, a *configuration* of Lib^k is a pair $\gamma = (d, \mu)$ where $d \in \mathbb{D}$ is the current value of the shared variable and μ is a map from $\{1, \ldots, k\}$ to $Q \uplus \{\bot\}$, specifying, for each thread i, the state in which the method called by thread i is. The symbol $\bot$ is used for threads which are *idle* (not calling any method at the moment).

A *step* from a configuration $\gamma = (d, \mu)$ to $\gamma' = (d', \mu')$ can be:

- thread i calling method j, denoted by $\gamma \xrightarrow{\texttt{call}(i, M_j)} \gamma'$, with $\mu(i) = \bot$, $\mu' = \mu[i \leftarrow q_0^j]$, and $d' = d$,
- thread i returning from method j, denoted by $\gamma \xrightarrow{\texttt{ret}(i)} \gamma'$, with $\mu(i) = q_f^j$, $\mu' = \mu[i \leftarrow \bot]$, and $d' = d$,
- thread i doing a read in method j, denoted by $\gamma \to \gamma'$ (no label) with $\mu(i) = q \in M_j$, $\mu' = \mu[i \leftarrow q']$, $(q, \mathsf{read}, d, q') \in \delta^j$, and $d' = d$,
- thread i doing a write in method j, denoted by $\gamma \to \gamma'$ (no label) with $\mu(i) = q \in M_j$, $\mu' = \mu[i \leftarrow q']$, $(q, \mathsf{write}, d', q') \in \delta^j$.

An *execution* of Lib^k is a sequence of steps $\gamma_0 \to \gamma_1 \ldots \to \gamma_l$ where $\gamma_0 = (d_0, \mu_0)$, with $\mu_0(i) = \bot$ for all i, is the initial configuration.

The *trace* h of an execution is the sequence of labels (call's and return's) of its steps. The set of traces of Lib^k is denoted by $Traces(Lib^k)$. Note that in a trace, a call event may be without a corresponding return event (if the method has not returned yet). In which case, the call event is said to be *open*. A trace with no open calls in called *complete*.

Given a complete trace h, we define for each pair of matching call and return events a *method event*. We say that a method event e_1 *happens before* another method event e_2 if the return event of e_1 is before the call event of e_2 in h; this defines a *happen-before* relation on the method events. The *label* of a method event is the method name corresponding to its call event.

2.2 Linearizability

Let h be a trace of $Traces(Lib^k)$ for some library Lib and integer k. A complete trace h' is said to be a *completion* of h if we can remove some (possibly zero) open calls from h, as well as close some others open calls (possibly zero) by adding return events at the end of h in order to obtain h'.

A *specification* for a library $Lib = \{M_1, \ldots, M_m\}$ is a language of finite words S over the alphabet $\{M_1, \ldots, M_m\}$.

Definition 1 (Linearizability). *A complete trace h is said to be* linearizable *with respect to a specification S if there exists a total order on the method events,*

respecting the happen-before order, such that the corresponding sequence of labels is a word in S. A trace h is said to be linearizable *with respect to S if it has a completion which is linearizable (with respect to S).*

Problem 1 (Linearizability). Input: A library $Lib = \{M_1, \ldots, M_m\}$, a non-deterministic finite automaton (NFA) S representing the specification, and an integer k given in unary.

Question: Are all the traces of $Traces(Lib^k)$ linearizable w.r.t. S?

Note: the size of the input is the size of all the automata appearing in the input (number of states + number of transitions + size of the alphabet) to which we add k.

We give in Figs 1, 2, and 3 some examples to illustrate Linearizability. To represente executions, we draw a method event as an interval, where the left end of the interval corresponds to the call event of the method event, and the right end corresponds to the return event. This way, when two method events overlap, they can be ordered arbitrarily, but when a method event e_1 is completely before a method event e_2, e_1 has to be ordered before e_2.

Above an interval, we write the name of the method corresponding to the method event, and below, we write the (unique) name of the method event.

For all the examples, we consider the regular language $S = (M_A M_B)^*$ as a specification. Fig 1 represents an execution which is linearizable, since its method events can be ordered as the sequence $e_1e_2e_3e_4$, whose corresponding sequence of labels is $M_AM_BM_AM_B$. Fig 2 represents an execution which is linearizable, since its method events can be ordered as the sequence $e_1e_2e_3e_4e_5e_6$, whose corresponding sequence of labels is $M_AM_BM_AM_BM_AM_B$. Fig 3 represents an execution which is similar to Fig 1 but is not linearizable.

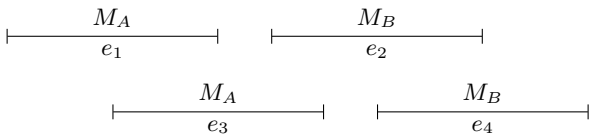

Fig. 1. A linearizable execution, which can be ordered as $e_1e_2e_3e_4$

2.3 Letter Insertion

We were able to define a new problem, *Letter Insertion*, which: 1) can be reduced to Linearizability, 2) is very easy to state (compared to Linearizability), 3) is still complex enough to capture the difficult part of Linearizability as we'll show it is EXPSPACE-hard.

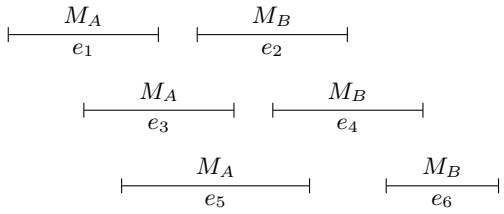

Fig. 2. A linearizable execution, which can be ordered as $e_1e_2e_3e_4e_5e_6$

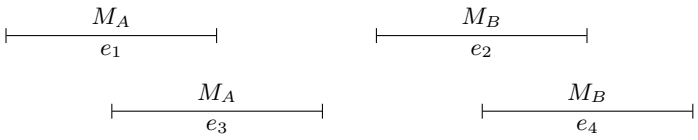

Fig. 3. A non-linearizable execution

Problem 2 (Letter Insertion). Input: A set of *insertable* letters $A = \{a_1, \ldots, a_l\}$. An NFA N over an alphabet $\Gamma \uplus A$.

Question: For all words $w \in \Gamma^*$, does there exist a decomposition $w = w_0 \cdots w_l$, and a permutation p of $\{1, \ldots, l\}$, such that $w_0 a_{p[1]} w_1 \ldots a_{p[l]} w_l$ is accepted by N?

Said differently, for any word of Γ^*, can we insert the letters $\{a_1, \ldots, a_l\}$ (each of them exactly once, in any order, anywhere in the word) to obtain a word accepted by N?

Note: the size of the input is the size of N, to which we add l.

3 Reduction from Letter Insertion to Linearizability

In this section, we show that Letter Insertion can be reduced in polynomial time to Linearizability. When we later show that Letter Insertion is EXPSPACE-hard, we will get that Linearizability is EXPSPACE-hard as well.

Intuitively, the letters $A = \{a_1, \ldots, a_l\}$ of Letter Insertion represent methods which are all overlapping with every other method, and the word w represents methods which are in sequence. Letter Insertion asks whether we can insert the letters in w in order to obtain a sequence of N while linearizability asks whether there is a way to order all the letters, while preserving the order of w, to obtain a sequence of N, which is equivalent.

Lemma 1. *Letter Insertion can be reduced in polynomial time to Linearizability.*

Proof. Let $A = \{a_1, \ldots, a_l\}$ and N an NFA over some alphabet $A \uplus \Gamma$.

Define k, the number of threads, to be $l + 2$.

We will define a library Lib composed of

- methods $M_1, \ldots, M_l$, one for each letter of A
- methods M_γ, one for each letter of Γ
- a method M_{Tick}.

and a specification S_N, such that (A, N) is a valid instance of Letter Insertion if and only if Lib^k is linearizable with respect to S_N.

For the domain of the shared variable, we only need three values: $\mathbb{D} = \{\mathsf{Begin}, \mathsf{Run}, \mathsf{End}\}$ with Begin being the initial value.

The methods M_γ are all identical. They just read the value Run from the shared variable (see Fig 4).

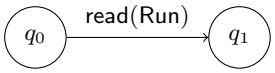

Fig. 4. Description of M_γ, $\gamma \in \Gamma$

The methods $M_1, \ldots, M_l$ all read Begin, and then read End (see Fig 5).

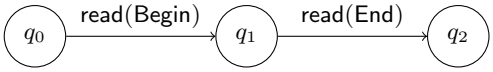

Fig. 5. Description of $M_1, \ldots, M_l$

The method M_{Tick} writes Run, and then End (see Fig 6).

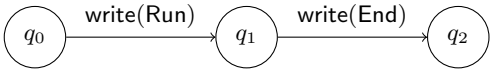

Fig. 6. Description of M_{Tick}

The specification S_N is defined as the set of words w over the alphabet $\{M_1, \ldots, M_l\} \cup \{M_{\mathsf{Tick}}\} \cup \{M_\gamma | \gamma \in \Gamma\}$ such that one the following condition holds:

- w contains 0 letter M_{Tick}, or more than 1, or
- for a letter M_i, $i \in \{1, \ldots, l\}$, w contains 0 such letter, or more than 1, or

– when projecting over the letters M_γ, $\gamma \in \Gamma$ and M_i, $i \in \{1, \ldots, l\}$, w is in N_M, where N_M is N where each letter γ is replaced by the letter M_γ, and where each letter a_i is replaced by the letter M_i.

Since N is an NFA, S_N is also an NFA. Moreover, its size is polynomial is the size of N. We can now show the following equivalence:

1. there exists a word w in Γ^*, such that there is no way to insert the letters from A in order to obtain a word accepted by N
2. there exists an execution of Lib with k threads which is not linearizable w.r.t. S_N

(1) $\implies$ (2). Let $w \in \Gamma^*$ such that there is no way to insert the letters A in order to obtain a word accepted by N. We construct an execution of Lib following Fig 7, which is indeed a valid execution.

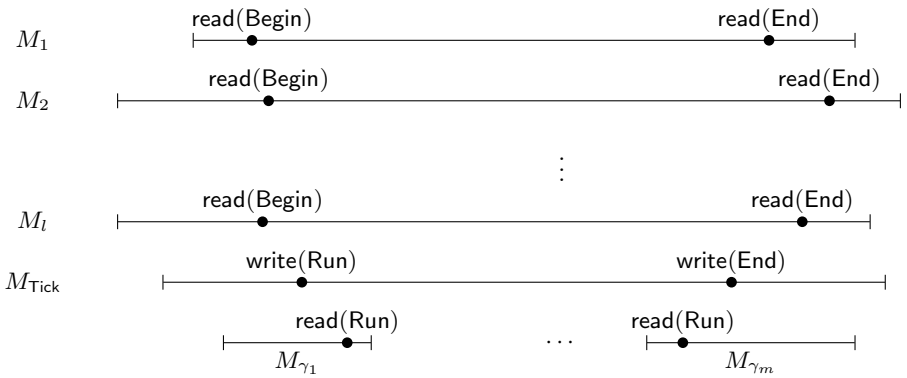

Fig. 7. Non-linearizable execution corresponding to a word $\gamma_1 \ldots \gamma_m$ in which we cannot insert the letters from $A = \{a_1, \ldots, a_l\}$ to make it accepted by N. The points represent steps in the automata.

This execution is not linearizable since

– it has exactly one M_{Tick} method, and
– for each $i \in \{1, \ldots, l\}$, it has exactly one M_i method, and
– no linearization of this execution can be in N_M, since there is no way to insert the letters A into w to be accepted by N.

Note: The value of the shared variable is initialized to Begin, allowing the methods M_i ($i \in \{1, \ldots, l\}$) to make their first transition. M_{Tick} then sets the value to Run, thus allowing the methods $M_\gamma, \gamma \in \Gamma$ to execute. Finally, M_{Tick} sets the value to End, allowing the methods $M_\gamma, \gamma \in \Gamma$ to make their second transition and return. This tight interaction will enable us to show in the second part of the proof that all non-linearizable executions of this library have this very particular form.

(2) $\implies$ (1). Let r be an execution which is not linearizable w.r.t. S_N. We first show that this execution should roughly be of the form shown in Fig 7. First, since it is not linearizable w.r.t. S_N, it must have at least one completed M_{Tick} method event. It it only had open M_{Tick} events (or no M_{Tick} events at all), it could be linearized by dropping all the open calls to M_{Tick}. Moreover, it cannot have more than M_{Tick} method event (completed or open), as it could also linearized, since S_N accepts all words with more than one M_{Tick} letter.

We can show similarly that for each $i \in \{1, \ldots, l\}$, it has exactly one M_i method which is completed (and none open).

Moreover, the methods M_i ($i \in \{1, \ldots, l\}$ can only start when the value of the shared variable is Begin, and they can only return after reading the value End. Since this value can only be changed (once) by the single M_{Tick} method of our executions, this ensures that the methods M_i ($i \in \{1, \ldots, l\}$) (and M_{Tick} itself) all overlap with one another, and with every other completed method.

This implies that the completed methods $M_\gamma, \gamma \in \Gamma$ can only appear in a single thread t (since $M_1, \ldots, M_l, M_{\mathsf{Tick}}$ already occupy $l+1$ threads amongst the $l+2$ available). Thus, we define $w \in \Gamma^*$ to be the word corresponding to the completed methods M_γ, $\gamma \in \Gamma$ of the execution in the order in which they appear in thread t.

Since r is not linearizable, we cannot insert M_i ($i \in \{1, \ldots, l\}$) into the completed methods of thread t order to be accepted by S_N. In particular, this implies that there is no way to insert the letters A in w in order to be accepted by N.

4 Letter Insertion is EXPSPACE-hard

We now reduce, in polynomial time, arbitrary exponentially bounded Turing machines, to the Letter Insertion problem, which shows it is EXPSPACE-hard. We first give a few notations.

A *deterministic Turing machine* $\mathcal{M}$ is a tuple (Q, δ, q_0, q_f) where:

- Q is the set of states,
- $\delta : (Q \times \{0,1\}) \to (Q \times \{0,1\} \times \{\leftarrow, \rightarrow\})$ is the transition function
- q_0, q_f are the initial and final states, respectively.

A computation of $\mathcal{M}$ is said to be *accepting* if it ends in q_f.

For the rest of the paper, we fix a Turing machine $\mathcal{M}$ and a polynomial P such that all runs of $\mathcal{M}$ starting with an input of size n use at most $2^{P(n)}$ cells, and such that the following problem is EXPSPACE-complete.

Problem 3 (Reachability). Input: A finite word t.

Question: Is the computation of $\mathcal{M}$ starting in state q_0, with the tape initialized with t, accepting?

Lemma 2 (Letter Insertion). *Letter Insertion is* EXPSPACE*-hard.*

Note: the sublemmas 3,4,5,6,7 are all part of the proof of Lemma 2.

Proof. We reduce in polynomial time the Reachability problem for EXPSPACE Turing machines to the *negation* of Letter Insertion. This still shows that Letter Insertion is EXPSPACE-hard, as the EXPSPACE complexity class is closed under complement.

Let t be a word of size n. Our goal is to define a set of letters A and an NFA N over an alphabet $\Gamma \uplus A$, such that the following two statements are equivalent:

- the run of $\mathcal{M}$ starting in state q_0 with the tape initialized with t is accepting (which, by definition of $\mathcal{M}$, uses at most $2^{P(n)}$ cells),
- there exists a word w in Γ^*, such that there is no way to insert (see Problem 2) the letters A in order to obtain a word accepted by N.

More specifically, we will encode runs of our Turing machine as words, and the automaton N, with the additional set of insertable letters A, will be used in order to detect words which:

- don't represent *well-formed sequence of configurations* (defined below),
- or represent a sequence of configurations where the initial configuration is not initialized with t and state q_0, or where the final configuration isn't in state q_f,
- or contain an error in the computation, according to the transition rules of $\mathcal{M}$.

A *configuration* of $\mathcal{M}$ is an ordered sequence $(c_0, \ldots, (q, c_i), \ldots, c_{2^{P(n)}-1})$ representing that the content of the tape is $c_0, \ldots, c_{2^{P(n)}-1} \in \{0,1\}$, the current control state is $q \in Q$, and the head is on cell i.

We denote by $\mathbf{i}$ the binary representation of $0 \leq i < 2^{P(n)}$ using $P(n)$ digits. Given a configuration, we represent cell i by: "$\mathbf{i} : c_i;$" if the head of $\mathcal{M}$ is not on cell i, and by "$\mathbf{i} : qc_i;$" if the head is on cell i and the current state of $\mathcal{M}$ is q. The configuration given above is represented by the word:

$$\$\mathbf{0} : c_0; \mathbf{1} : c_1; \ldots \mathbf{i} : qc_i; \ldots \mathbf{2^P(n) - 1} : c_{2^{P(n)}-1}; \hookleftarrow$$

.

Words which are of this form for some $c_0, \ldots, c_{2^{P(n)}-1} \in \{0,1\}$, $q \in Q$, are called *well-formed configurations*. A sequence of configurations is then encoded as $\triangleright \mathsf{cfg_1} \ldots \mathsf{cfg_k} \square$ where each $\mathsf{cfg_i}$ is a well-formed configuration. A word of this form is called a *well-formed sequence of configurations*. We now fix Γ to be $\{0, 1, \triangleright, \square, \$, \hookleftarrow, ;, :\}$.

Lemma 3. *There exists an NFA N_{notWF} of size polynomial in n, which recognizes words which are not well-formed configurations.*

Proof. A word is not a well-formed configuration if and only if one of the following holds:

- it is not of the form $\$((0+1)^{P(n)} : (Q+\epsilon)(0+1);)^* \hookleftarrow$, or

– it has no symbol from Q, or more than one, or
– it doesn't start with $\$\mathbf{0}:$, or
– it doesn't end with $\mathbf{2^{P(n)}-1} : (Q+\epsilon)(0+1); \hookleftarrow$, or
– it contains a pattern $\mathbf{i} : (Q+\epsilon)(0+1); \mathbf{j}$: where $j \neq i+1$.

For all violations, we can make an NFA of size polynomial in n recognizing them, and then take their union. The most difficult one is the last, for which there are detailed constructions in Fürer [6] and Mayer and Stockmeyer [10].

We here give a sketch of the construction. Remember that **i** and **j** are binary representation using $P(n)$ bits. We want an automaton recognizing the fact that $j \neq i+1$. The automaton guesses the least significant bit b ($P(n)$ possible choices) which makes the equality $i+1=j$ fails, as well as the presence or not of a carry (for the addition $i+1$) at that position. We denote by $\mathbf{i}[b]$ the bit b of **i** and likewise for **j**. Then, the automaton checks that: 1) there is indeed a violation at that position (for instance: no carry, $\mathbf{i}[b]=0$ and $\mathbf{j}[b]=1$) and 2) there is carry if and only if all bits less significant that b are set to 1 is **i**.

Lemma 4. *There exists an NFA $N_{\mathsf{NotSeqCfg}}$ of size polynomial in n, which recognizes words which:*

– *are not a well-formed sequence of configurations, or where*
– *the first configuration is not in state q_0, or*
– *the first configuration is not initialized with t, or*
– *the last configuration is not in state q_f.*

Proof. Non-deterministic union between N_{notWF} and simple automata recognizing the last three conditions.

The problem is now in making an NFA which detects violations in the computation with respect to the transition rules of $\mathcal{M}$. Indeed, in our encoding, the length of one configuration is about $2^{P(n)}$, and thus, violations of the transition rules from one configuration to the next are going to be separated by about $2^{P(n)}$ characters in the word. We conclude that we cannot make directly an automaton of polynomial size which recognize such violations.

This is where we use the set of insertable letters A. We are going to define and use it here, in order to detect words which encode a sequence of configurations where there is a computation error, according to the transition rules of $\mathcal{M}$.

The set A, containing $2P(n)$ new letters, is defined as $A = \{p_1, \ldots, p_{P(n)}, m_1, , \ldots, m_{P(n)}\}$.

We want to construct an NFA N_{NotDelta}, such that, for a word w which is a well-formed sequence of configurations, these statements are equivalent:

– w has a computation error according to the transition rules δ of $\mathcal{M}$
– we can insert the letters A in w to obtain a word accepted by N_{NotDelta}.

The idea is to use the letters A in order to identify two places in the word corresponding to the same cell of $\mathcal{M}$, but at two successive configurations of the run.

As an example, say we want to detect a violation of the transition $\delta(q, 0) = (q', 1, \rightarrow)$, that is, which reads a 0, writes a 1, moves the head to the right, and changes the state from q to q'.

Assume that w contains a sub-word of the following form:

$$\mathbf{i} : q0; \ldots \$ \ldots \mathbf{i} : 1; \mathbf{i} + \mathbf{1} : q'' c_{i+1};$$

where q'' is different than q'

The single \$ symbol on the middle of the sub-word ensures that we are checking violations in successive configurations. Here, with the current state being q, the head read 0 on cell i, wrote 1 successfully, and moved to the right. But the state changed to q'' instead of q'. Since we assumed that $\mathcal{M}$ is deterministic, this is indeed a violation of the transition rules.

We now have all the ingredients in order to construct N_{NotDelta}. It will be built as a non-deterministic choice (or union) of N_t for all possible transitions $t \in \delta$ (with δ seen as a relation).

As an example, we show how to construct the automaton $N^{(1)}_{((q,0),(q',1,\rightarrow))}$, part of N_{NotDelta}, and recognizing violations of $\delta(q, 0) = (q', 1, \rightarrow)$, where the head was indeed moved to right, but the state was changed to some state q'' instead of q', like above. Other violations may be recognized similarly.

$N^{(1)}_{((q,0),(q',1,\rightarrow))}$ starts by finding a sub-word of the form (the + denotes the disjunction or union of regular expressions, and * denotes the Kleene star, 0 or more repetitions):

$$(m_1 0 + p_1 1) \ldots (m_{P(n)} 0 + p_{P(n)} 1) : q0; \tag{1}$$

meaning the state is q and the head points to a cell containing 0. After that, it reads arbitrarily many symbols, but exactly one \$ symbol, which ensures that the next letters it reads are from the next configuration. Finally, it looks for a sub-word of the form

$$(p_1 0 + m_1 1) \ldots (p_{P(n)} 0 + m_{P(n)} 1) : (0 + 1); (0 + 1)^* : q;'' \tag{2}$$

for some $q'' \neq q'$.

We can now show the following.

Lemma 5. *For a well-formed sequence of configurations w, these two statements are equivalent:*

1. *there is a way to insert the letters A into w to be accepted by $N^{(1)}_{((q,0),(q',1,\rightarrow))}$*
2. *in the sequence of configurations encoded by w, there is a configuration where the state was q and the head was pointing to a cell containing 0, and in the next configuration, the head was moved to the right, but the state was not changed to q' (computation error).*

Proof. ($\Leftarrow$). We insert the letters A in front of the binary representation of the cell number where the violation occurs. The violation involves two configurations:

in the first, we insert m's in front of 0's, and p's in front of 1's, and in the second, it's the other way around.

This way, we inserted all the letters of A (exactly) once into w, and $N^{(1)}_{((q,0),(q',1,\rightarrow))}$ is now able to recognize the patterns (1) and (2) described above.

($\Rightarrow$). For the other direction, let w be a well-formed sequence of configurations such that there exists a way to insert the letters A into w, in order to obtain a word w_A accepted by $N^{(1)}_{((q,0),(q',1,\rightarrow))}$.

Since each letter of A can be inserted only once, the sub-word matched by $(m_1 0 + p_1 1) \dots (m_{P(n)} 0 + p_{P(n)} 1)$ in pattern (1) in $N^{(1)}_{((q,0),(q',1,\rightarrow))}$ has to be the same as the one matched by $(p_1 0 + m_1 1) \dots (p_{P(n)} 0 + m_{P(n)} 1)$ in pattern (2), up to exchanging m's and p''s.

Moreover, having exactly one \$ symbol in between the two patterns ensures that they correspond to the same cell, but in two successive configurations.

Finally, the facts that q'' is different that q' and that $\mathcal{M}$ is deterministic ensures that the sequence of configurations represented by w indeed contains a computation error according to the rule $\delta(q, 0) = (q', 1, \rightarrow)$.

We thus get the following lemma for the automaton N_{NotDelta}.

Lemma 6. *For a word w which is well-formed sequence of configurations, these statements are equivalent:*

- *we can insert the letters A in w to obtain a word accepted by N_{NotDelta},*
- *w has a computation error according to the transition rules δ of $\mathcal{M}$.*

Proof. Construct all the N_t for $t \in \delta$ (with δ considered as a relation). Construct similarly an automaton recognizing the violation where a cell changes while the head was not here. Take the union of all these automata, the proof then follows from Lemma 5.

By taking the union $N = N_{\mathsf{NotSeqCfg}} \cup N_{\mathsf{NotDelta}}$, we finally get the intended result, which ends the reduction.

Lemma 7. *The following two statements are equivalent.*

- *the run of $\mathcal{M}$ starting in state q_0 with the tape initialized with t is accepting,*
- *there exists a word w in Γ^*, such that there is no way to insert the letters A in order to obtain a word accepted by N.*

Proof. ($\Rightarrow$) Let w be the well-formed sequence of configurations representing the sequence of configurations of the accepting run in $\mathcal{M}$, with the tape initialized with t. Then by Lemma 4 and Lemma 6, there is no way to insert the letters A in order to obtain a word accepted by $N_{\mathsf{NotSeqCfg}}$ or N_{NotDelta}.

($\Leftarrow$) Let $w \in \Gamma^*$ be a word such that there is no way to insert the letters A in order to obtain a word accepted by N. First, since w is not accepted by $N_{\mathsf{NotSeqCfg}}$, it represents a well-formed sequence of configurations, starting in state q_0 with the tape initialized with t and ending in state q_f (Lemma 4). Moreover, since there is no way to insert the letters to obtain a word from N_{NotDelta}, w has no computation error according to the transition rules δ of $\mathcal{M}$ (Lemma 6).

This ends the proof of Lemma 2.

Since Letter Insertion is EXPSPACE-hard and, Letter Insertion reduces to Linearizability, we get the main result of the paper.

Theorem 1 (Linearizability). *Linearizability is* EXPSPACE*-complete.*

Proof. It was previously shown that Linearizability is in EXPSPACE [1]. EXPSPACE-hardness follows from Lemmas 1 and 2

5 Conclusion

We define a new problem, Letter Insertion, simpler than Linearizability, but still hard enough to capture the main difficulties of Linearizability. We showed that the Letter Insertion problem is EXPSPACE-hard, and could thus deduce that the Linearizability problem is EXPSPACE-hard.

Our result applies even with all the following restrictions: the number of threads is given in unary, there is a unique shared variable whose domain size is 3, the library has a constant number of automata "shapes" (3 in our reduction) using less than 3 states, the methods of the library are deterministic, the methods of the library have no loop, and the instructions within the methods can only read or write, but never do both atomically.

For future work, we plan to show that restricting ourselves to deterministic specifications (using a DFA instead of an NFA in the input of the problem) does not reduce the complexity. Furthermore, it would be interesting to find a large class of specifications including the most common ones (stack, queue, ...) for which our lower-bound does not apply and where we could reduce the complexity.

Bibliography

[1] Alur, R., McMillan, K.L., Peled, D.: Model-checking of correctness conditions for concurrent objects. Inf. Comput. 160(1-2), 167–188 (2000)

[2] Bouajjani, A., Emmi, M., Enea, C., Hamza, J.: Verifying concurrent programs against sequential specifications. In: ESOP '13. LNCS, vol. 7792, pp. 290–309. Springer (2013)

[3] Burckhardt, S., Dern, C., Musuvathi, M., Tan, R.: Line-up: A complete and automatic linearizability checker. In: Proceedings of the 2010 ACM SIGPLAN Conference on Programming Language Design and Implementation. pp. 330–340. PLDI '10, ACM, New York, NY, USA (2010), http://doi.acm.org/10.1145/1806596.1806634

[4] Elmas, T., Qadeer, S., Sezgin, A., Subasi, O., Taşıran, S.: Simplifying linearizability proofs with reduction and abstraction. In: TACAS '10: Proc. 16th Intl. Conf. on Tools and Algorithms for the Construction and Analysis of Systems. LNCS, vol. 6015, pp. 296–311. Springer (2010)

[5] Filipovic, I., O'Hearn, P.W., Rinetzky, N., Yang, H.: Abstraction for concurrent objects. Theor. Comput. Sci. 411(51-52), 4379–4398 (2010)

[6] Fürer, M.: The complexity of the inequivalence problem for regular expressions with intersection. In: Proceedings of the 7th Colloquium on Automata, Languages and Programming. pp. 234–245. Springer-Verlag, London, UK, UK (1980), http://dl.acm.org/citation.cfm?id=646234.682559

[7] Gibbons, P.B., Korach, E.: Testing shared memories. SIAM J. Comput. 26(4), 1208–1244 (1997)

[8] Herlihy, M., Wing, J.M.: Linearizability: A correctness condition for concurrent objects. ACM Trans. Program. Lang. Syst. 12(3), 463–492 (1990)

[9] Hunt, H.: The Equivalence Problem for Regular Expressions with Intersection is Not Polynomial in Tape. Department of Computer Science: Technical report, Cornell University, Department of Computer Science (1973), http://books.google.fr/books?id=52j6HAAACAAJ

[10] Mayer, A.J., Stockmeyer, L.J.: The complexity of word problems - this time with interleaving. Inf. Comput. 115(2), 293–311 (Dec 1994), http://dx.doi.org/10.1006/inco.1994.1098

[11] Rajamani, S.K., Walker, D. (eds.): Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015. ACM (2015), http://dl.acm.org/citation.cfm?id=2676726

[12] Vafeiadis, V.: Automatically proving linearizability. In: CAV '10. LNCS, vol. 6174, pp. 450–464 (2010)

[13] Vechev, M.T., Yahav, E., Yorsh, G.: Experience with model checking linearizability. In: SPIN '09: Proc. 16th Intl. SPIN Workshop on Model Checking Software. LNCS, vol. 5578, pp. 261–278. Springer (2009)